MAC-Assisted Broadcast Speedup in Ad-Hoc Wireless Networks

Ashikur Rahman and Pawel Gburzynski
Dept. of Computing Science, University of Alberta
Edmonton, Alberta, Canada
ashikur@cs.ualberta.ca, pawel@cs.ualberta.ca

ABSTRACT

The primary performance objective of a broadcast scheme in an ad-hoc wireless network is to reduce the total number of retransmissions needed to reach all nodes. Another (less appreciated) measure of interest is the broadcast latency, i.e., the amount of time required to complete the operation. We point out the tradeoff between the two measures and show how to significantly reduce the broadcast latency in networks whose MAC schemes are derivatives of IEEE 802.11.

Categories and Subject Descriptors: C.2.2 [Network Protocols]: Routing protocols

General Terms: Algorithms, Performance.

Keywords: Ad-hoc Networks, Broadcast, Medium Access Control.

1. INTRODUCTION

Broadcasting in ad-hoc wireless networks is a fundamental and frequently used operation. Reactive routing protocols such as DSR [4], AODV [12], TORA [8], depend on broadcasting for route discovery. Proactive protocols, such as DSDV [10] and WRP [6], periodically broadcast updated information about cost metrics. Some protocols even use (selective) broadcasting for the actual forwarding [14].

The common goal of all algorithms proposed in the literature is minimizing the number of retransmissions while ensuring that the broadcast message reaches all nodes in the network. To the best of our knowledge, none of those efforts explicitly addresses the issue of broadcast latency. For example, in a reactive routing protocol, the broadcast latency directly determines how quickly routes are discovered. For a proactive protocol, which disseminates the cost metric information via broadcasting, the high cost of this operation will reduce the optimum update frequency and render small updates too costly to disseminate.

We propose a modification to the backoff component of the collision avoidance schemes derived from IEEE 802.11 that significantly reduces the broadcast latency while keeping the number of retransmissions small. Our approach is somewhat reminiscent of the previous work [2] on Quality of Service issues related to fairness and priority scheduling. In our own previous work [13, 14], we proposed two other modifications to IEEE 802.11 to facilitate topology control and increase the reliability of broadcast-based forwarding.

IWCMC'06, July 3–6, 2006, Vancouver, British Columbia, Canada.

2. THE OPTIMIZATION FRAMEWORK

As shown in [5], the problem of constructing a minimum flooding tree (the so-called *Minimum Connected Dominating Set*), which would identify the best collection of nodes to retransmit a brodcast packet as to reach all other nodes, is NP-hard. Thus, most approaches to efficient broadcasting in ad-hoc networks start with flooding and attempt to reduce its extent with the assistance of some heuristics. Typically, such schemes impose delays before retransmissions. While waiting, a node monitors the traffic in the neighborhood, which may let it conclude that its own retransmission would be redundant. This generic algorithm (which we shall refer to as GDA) is shown in Figure 1. The parts in bold, i.e., the setting of the defer timer and the pruning criteria are the specific features of any actual scheme.

1. *Packet reception: if the packet was seen before, drop it (duplicate discard); otherwise, push it into a queue.*
2. **Set the defer timer to some value** D *and wait.*
3. *Keep receiving broadcast packets from neighbors. When the timer goes off, proceed to 4.*
4. *End of waiting:* ***evaluate the pruning criteria.*** *If the packet appears to be redundant, drop it; otherwise, pass it to the MAC layer for retransmission.*

Figure 1: GBA: the Generic Broadcast Algorithm

The amount of defer time D determines the size of the window of opportunity for the acquisition of input to the pruning criteria. One obvious idea is to set the timer to some random value [7] (note that setting it to a constant value, the same for all nodes, makes little sense). Another approach [16] uses a distance-based formula that favors retransmissions by distant nodes.

The popular pruning criteria [7] include the naive *probabilistic approach*, whereby the packet is rebroadcast with probability P and ignored with probability $(1 - P)$, the *counter-based approach* (the packet is considered redundant if more than N of its copies have been received from the neighbors during the waiting time D), the *distance-based approach* (the packet is considered redundant if the node has received another copy of the packet transmitted by a neighbor located not further than some distance d from the node), and the *location-based approach* (the packet is redundant if the node falls into the convex hull formed by those of its neighbors that have retransmitted the packet during the waiting interval D). More refined criteria involve neighbor information passed in the headers of broadcast packets, e.g., one-hop neighbors [5], two-hop neighbors [9], or generally k-hop neighbors [15, 17]. The packet is considered redundant if the combined list of nodes extracted from the headers of all its copies received by the node within the interval D covers the node's entire neighborhood. It seems to us that pushing this approach beyond the two-hop case, assumed in SBA (the *Ssalable Broadcast Algorithm* [9]), results in diminishing returns.

Except for the most naive probabilistic criterion, it is natural to expect that the longer the interval D, the more information is likely to reach the node before it is forced to make the decision. This way, longer latency will result in fewer retransmissions and vice-versa. Note that in all cases it makes sense to randomize D. Otherwise, one-hop neighbors receiving the same packet at the same time will never exchange any feedback among themselves. Also note that only the criteria based on neighborhood information [5, 9, 15, 17] guarantee the reachability of all nodes.

Suppose a node u transmits a broadcast packet, which will be received by all it neighbors denoted by $\mathcal{N}(u)$. Each node $v \in \mathcal{N}(u)$ will schedule its retransmission after some delay D_u. Not all of those retransmissions are equally effective and important. For example, if v is located close to u, it can only reach a small additional area compared to what has been already covered by u. On the other hand, if v is located near the perimeter of the u's transmission range, it can relay the packet to more distant regions; thus, its retransmission is likely to cover many additional nodes. Consequently, it makes sense to give priority to distant nodes. In Figure 1, there appears to be a single place when the node can affect the timing of its retransmission, i.e., step 2. However, in reality, there is another "defer time" contributed to the packet's retransmission delay by the IEEE 802.11 MAC layer [3]. Under contention, a node has to wait for a certain number of idle slots chosen at random in the range of $[0, cw-1]$, where cw is the so-called *contention window*. Thus, the total delay until retransmission is $D_{total} = D + D_{MAC}$, where D_{MAC} is the component incurred by the contention resolution algorithm of the MAC layer. Note that even if the node attempts to follow a sensible priority scheme in setting D, its intentions can be thwarted by the MAC layer, i.e., the actual retransmission order may be different from the intended one.

One could always choose $D >> D_{MAC}$, so that the range of the MAC-incurred component is insignificant. This way, however, the proper ordering will be achieved at the cost of increased latency. A better way would be to include the MAC delay in the overall scheme. This will bring about two benefits: (i) the nodes will be able to prioritize their retransmissions without the confusing impact of the MAC layer; (ii) the MAC layer will be able to use longer delays without having to worry about impairing the network layer. Thus, it may do a better job as far as avoiding collisions is concerned.

3. MAC-ASSISTED BROADCASTING

Our goal is to incorporate priorities into the calculation of D, which represents the complete defer delay, including the MAC component. In fact, we do not want to separate the two components at all. Whenever a node receives a broadcast packet, it calculates its priority and immediately passes the packet, along with its priority, to the MAC layer. The network layer reserves the right to revoke the retransmission request if, while the packet is awaiting its turn, the node concludes that the packet is redundant. The MAC layer will push the packet into the interface queue and initiate the backoff procedure. The way the backoff delay is calculated will be biased, such that high-priority packets will be delayed for a shorter time than low-priority packets.

Suppose that the retransmission priority has been quantized into n discrete levels, e.g., $1, 2, \ldots, n$, where n is a preconfigured parameter, and higher values indicate a higher priority. Assume that the priority of a given packet is i. The complete (backoff) delay of the packet is prescribed by:

$$R = (n-i) \times \frac{cw}{2^{\lceil log_2 n \rceil}} + \lfloor U(0,1) \times \frac{cw}{2^{\lceil log_2 n \rceil}} \rfloor \ , \qquad (1)$$

where $U(0,1)$ denotes uniform distribution between 0 and 1.

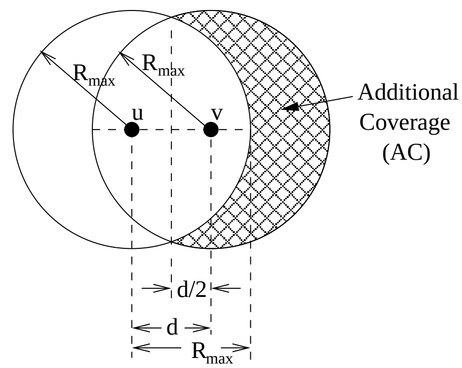

Figure 2: Additional coverage area

Assume that a node v has received a broadcast packet from node u (see Figure 2). Let d_{uv} be the distance between u and v, and the R_{max} stand for the maximum communication range. Then the additional coverage (denoted by AC) is expressed by:

$$\text{AC} = 2R_{max}^2 sin^{-1} \frac{d_{uv}}{2R_{max}} + \frac{d_{uv}}{2}\sqrt{4R_{max}^2 - d_{uv}^2} \ .$$

In accord with the quantization of priorities, it makes sense to divide the circle representing the coverage of u into n equal-width partitions. Specifically, node v is said to fall into partition i, $1 \le i \le n$ iff,

$$\frac{R_{max} \times (i-1)}{n} < d_{sv} \le \frac{R_{max} \times i}{n} \ .$$

Note that each partition corresponds to a circular stripe of width R_{max}/n. Node v will assign priority i to a broadcast packet received from u.

Consider the three nodes shown in Figure 3. Both nodes v and w are located within the transmission range of u, i.e., $d_{uv} < R_{max}$ and $d_{uw} < R_{max}$; however, $d_{uv} < R_{max}/2$ while $d_{uw} > R_{max}/2$. Let us assume that the counter

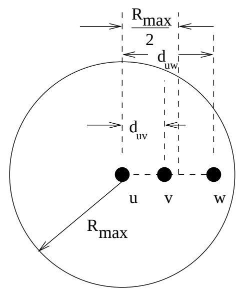

Figure 3: A 3-node scenario

threshold $N = 1$, i.e., if the same broadcast packet is received twice, the node will not retransmit the packet. Suppose that u transmits a packet, which is received by both v and w. The two nodes calculate their delays D_v and D_w and schedule their retransmissions.

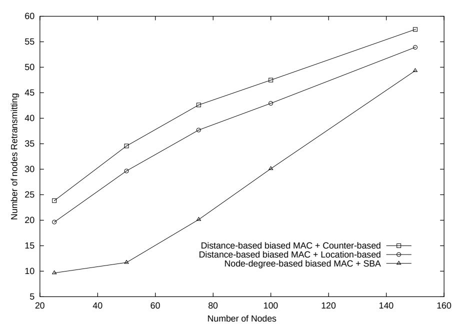

(a)

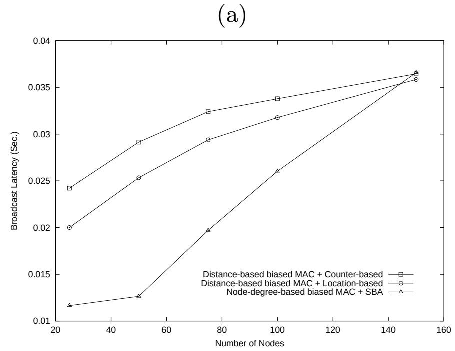

(b)

Figure 4: Performance of biased MAC combined with different pruning schemes w.r.t.: (a) the number of retransmissions, (b) the latency.

If $D_v < D_w$, then node v will go first. Its transmission will make the counter at w reach 2 before D_w expires, and, consequently, w will not transmit. On the other hand, if $D_v > D_w$, then w will retransmit and v will not. One can easily see that as long as the priority scheme used in determining D_v and D_w follows Formula 1, and $d_{uw} - d_{uv} > R_{max}/n$, we will have the second scenario, in agreement with the principle of preferring nodes with larger additional coverage.

The location-based algorithm described in [7] can also benefit from the distance-based priority scheme. In this case, the role of the distance-based priority is to speed up

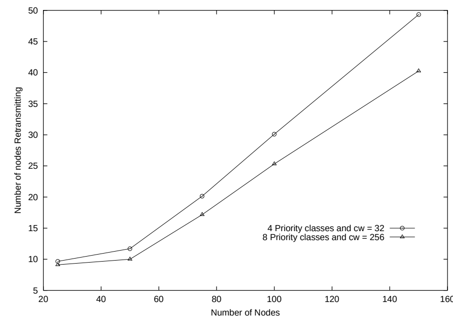

(a)

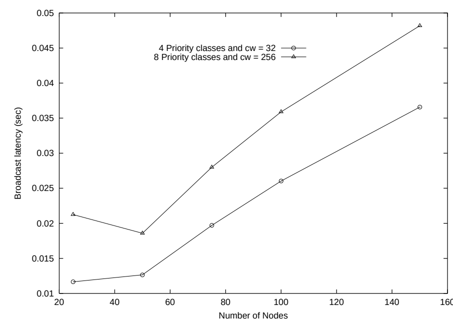

(b)

Figure 5: Performance of biased MAC with different numbers of priority classes and initial contention window size w.r.t.: (a) the number of retransmissions, (b) the latency.

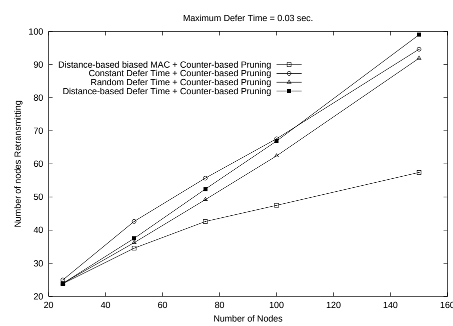

(a)

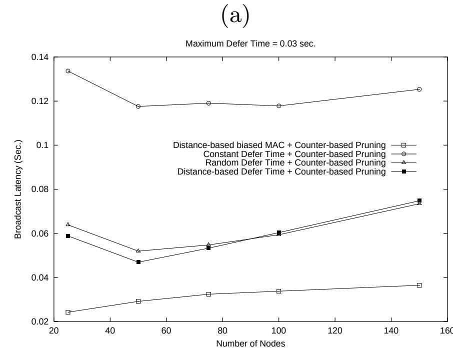

(b)

Figure 6: Comparison of different defer schemes combined with counter-based pruning w.r.t: (a) the number of retransmissions, (b) the latency.

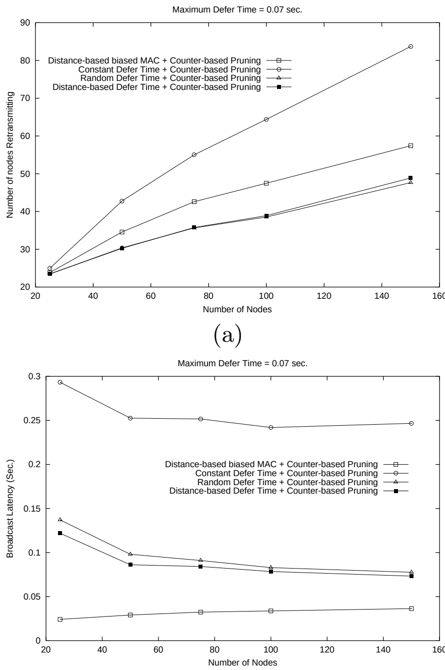

(b)

Figure 7: Comparison of different defer schemes combined with counter-based pruning w.r.t.: (a) the number of retransmissions, (b) the latency.

the transmission by distant nodes, which are more likely to cover more area and preempt the nodes located closer to the packet source. Even if the unassisted variant of the protocol could eventually find out which nodes should transmit and which should not, by giving preference to distant nodes, the consensus can be reached quicker, which will have a positive impact on the latency. Another way of prioritizing retransmissions by adjusting the defer time is to take into account the network density in the node's neighborhood. The general idea is to support retransmissions by nodes with many neighbors. This approach will locally maximize the rate at which broadcast packets spread through the network. More specifically, a node u receiving a broadcast packet for a possible retransmission will determine the value of the following three parameters: g_u–its own node degree (the number of its neighbors), g_{max}–the maximum node degree among its neighbors (including itself), g_{min}–the minimum node degree among its neighbors (including itself). Let $g_{var} = g_{max} - g_{min} + 1$ be a measure of variation in the node degree in u's neighborhood. The node will set its retransmission priority to i such that:

$$\frac{g_{var} \times (i-1)}{n} + g_{min} \leq g_u < \frac{g_{var} \times i}{n} + g_{min}$$

where $1 \leq i \leq n$ and n, as before, is the assumed number of discrete priority levels. This formula gives higher priority to nodes having more neighbors.

Degree-based priority schemes, like the one suggested above, will tend to reduce the latency of algorithms in the SBA class [9]. SBA assumes that every node knows the identities of its 2-hop neighbors. Having identified the last sender u of a received broadcast packet, node v can easily determine which of its neighbors are included in the neighbor set of u. As it receives more copies of the packet, v can find out whether those packets have completely covered its own set of neighbors. Should that happen before the defer time expires, v will cancel the scheduled retransmission. It is clear that by giving a head start to the nodes with large sets of neighbors, we will speed up the distribution of the broadcast packet in the neighborhood and faster prune out the redundant retransmissions.

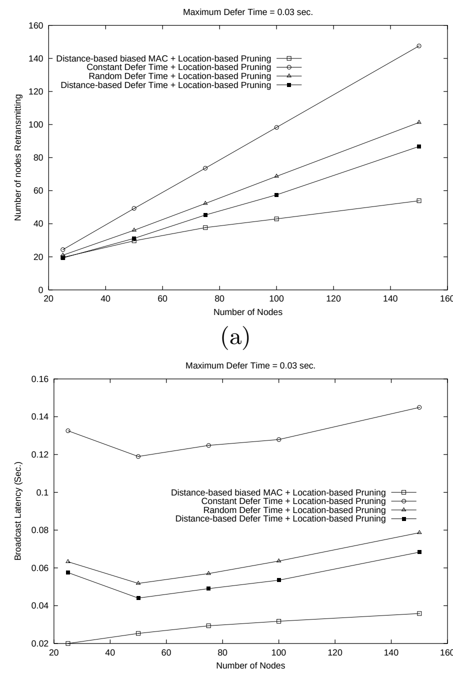

(b)

Figure 8: Comparison of different defer schemes combined with location-based pruning w.r.t: (a) the number of retransmitting nodes, (b) broadcast latency.

4. EXPERIMENTAL RESULTS

We built a detailed simulation model based on *ns*-2 [1] with wireless extensions. A modified version of the distributed coordination function (DCF) of the IEEE standard 802.11 [3], was used as the MAC layer. The modifications consisted in adding the capability to generate biased backoff delays, as described in Section 3.

We deployed networks with $25 - 150$ nodes uniformly distributed over a flat square area of $670m \times 670m$. The transmission range, $250m$, was the same for all nodes. Three pruning algorithms were implemented: a counter-based scheme with the threshold $N = 6$, a location-based scheme using the convex hull mechanism proposed in [7], and SBA.

Note that neither the counter-based nor the location-based schemes guarantee that every broadcast packet will always reach all nodes. The threshold values during our experiments were suitably chosen such that the reachability of the broadcast packets in all scenarios was more than 98%.

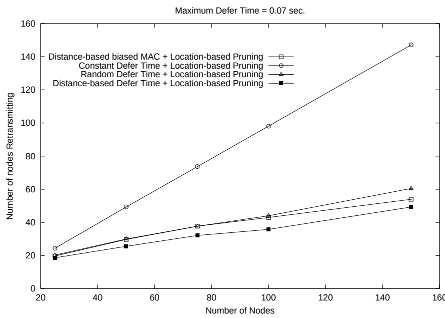

(a)

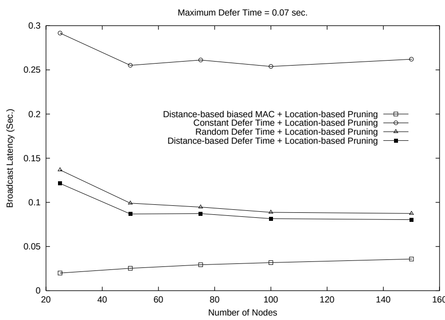

(b)

Figure 9: Comparison of different defer schemes combined with location-based pruning w.r.t.: (a) the number of retransmitting nodes, (b) broadcast latency.

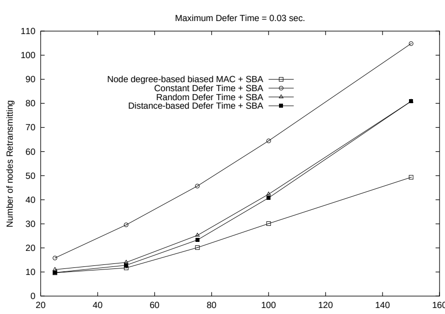

(a)

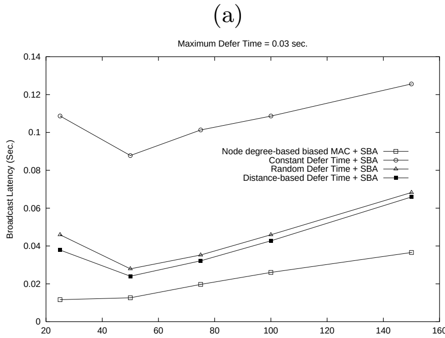

(b)

Figure 10: Comparison of different defer schemes combined with SBA w.r.t.: (a) the number of retransmissions, (b) the latency.

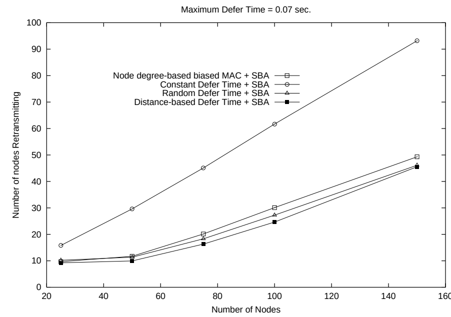

(a)

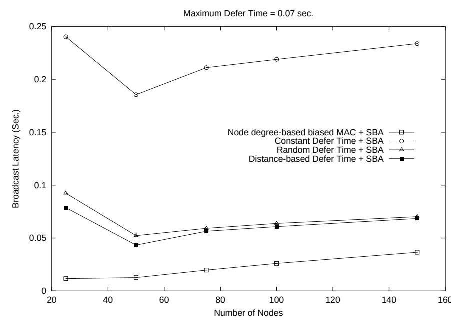

(b)

Figure 11: Comparison of different defer schemes combined with SBA w.r.t.: (a) the number of retransmissions, (b) the latency.

Each node in the network generated a broadcast packet in turn, but only one broadcast packet (possibly in multiple copies) was present in the network at any given time. The performance measures were collected for each broadcast and then, at the end of run, the averages of those measures were taken. The points used to make up the curves were the averages of those averages collected over five independent experiments.

Figure 4 illustrates the effect of MAC-biasing on different pruning strategies. Both distance-based and degree-based priority schemes were investigated, with four priority levels in each case.

Figure 4(a) shows the average number of nodes participating in retransmissions to complete the broadcast. The combination of the degree-based priority scheme with SBA pruning exhibits the best performance. With the distance-based priority scheme, the location-based pruning shows better performance than the counter-based approach.

Figure 4(b) shows the broadcast latency. Again the combination of the degree-based priority scheme and SBA exhibits the best performance. The location-based pruning also has a lower latency than the counter-based scheme with the same distance-based biased MAC.

According to the IEEE 802.11 standard, a node having a packet to transmit, starts with the contention window (cw) of 32 slots. Owing to the fact that network-layer delays needed by the priority schemes tend to be longer than typical MAC-layer delays (needed solely for contention resolution), it may make sense to start from a larger contention window. From the viewpoint of contention resolution, a larger window can only help. Thus, we ran a series of experiments to study the impact of a large initial contention window on the performance of our modifications.

With *cw* set to 256, we were able to accommodate more priority classes, so we increased their number from 4 to 8. The results (with SBA used as the pruning scheme) are shown in Figure 5. They indicate that the number of retransmissions can be reduced even further, but only at the cost of increased latency.

Figures 6–11 compare the performance of MAC-biasing with other defer schemes: constant, random, and distance-based—in combination with different pruning techniques. The standard 802.11 MAC protocol was used in the last three cases.

Our approach consistently offers a significantly lower latency than the other schemes, which, in some cases, comes at the price of a (slightly) increased number of retransmissions. However, the penalty in the number of retransmissions only occurs in those situations when the limit on defer time is high, thus encouraging longer delays during which the nodes have plenty of time to monitor activities in the neighborhood before deciding whether to retransmit. Even in the worst case (Figure 7(a)), the number of retransmissions is only about 20% worse compared to the best result (for random defer time), while the latency offered by our scheme is lower by the factor of two.

5. CONCLUSIONS

We have proposed a simple modification to the backoff algorithm of IEEE 802.11 aimed at improving the performance of broadcasting in wireless networks. Our enhancement considerably reduces the latency of broadcasting while maintaining a low number of retransmissions.

As a future work, we plan to investigate the performance of our scheme used as a supplement to some well-established reactive routing protocols, such as AODV [11] and DSR [4], to assist them in route discovery. It will be interesting to see how much the improvements to this part of the routing protocol will benefit the performance of the entire system. At the same time, we are modifying DSDV [10] to use our broadcasting technique for propagating routing tables.

6. REFERENCES

[1] The Network Simulator: NS-2: notes and documentation. http://www.isi.edu/nsnam/ns/.

[2] J. Deng and R.-S. Chang. A priority scheme for IEEE 802.11 DCF access method. *IEICE Trans. Commun.*, E82-B(1), January 1999.

[3] IEEE Standards Department. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 1997. IEEE standard 802.11-1997.

[4] D. B. Johnson and D. A. Maltz. Dynamic Source Routing in ad hoc wireless networks. In Imielinski and Korth, editors, *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.

[5] H. Lim and C. Kim. Multicast tree construction and flooding in wireless ad hoc networks. In *ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWIM)*, 2000.

[6] S. Murthy and J. J. Garcia-Luna-Aceves. An efficient routing protocol for wireless networks. *ACM Mobile Networks and Applications Journal*, pages 183–197, October 1996.

[7] Sze-Yao Ni, Yu-Chee Tseng, Yuh-Shyan C hen, and Jang-Ping Sheu. The broadcast storm problem in a mobile ad hoc network. In *Mobicom*, 1999.

[8] V.D. Park and M.S. Cors, on. A performance comparison of TORA and ideal link state routing. In *Proceedings of IEEE Symposium on Computers and Communications '98*, June 1998.

[9] W. Peng and X. Lu. On the reduction of broadcast redundancy in mobile ad hoc networks. In *Mobihoc*, 2000.

[10] C.E. Perkins and P. Bhagwat. Highly dynamic Destination-Sequenced Distance Vector routing (DSDV) for mobile computers. In *Proceedings of SIGCOMM'94*, pages 234–244, August 1993.

[11] C.E. Perkins and E.M. Royer. Ad-hoc On-demand Distance Vector Routing (AODV). In *Proceedings of the IEEE workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 90–100, 1999.

[12] C.E. Perkins, E.M. Royers, and S.R. Das. Ad-hoc On-demand Distance Vector Routing (AODV), February 2003. Internet Draft: draft-ietf-manet-aodv-13.txt.

[13] A. Rahman and P. Gburzynski. On constructing minimum-energy path-preserving graphs for ad hoc wireless networks. In *Proceedings of ICC '05*, Seoul, Korea, May 2005.

[14] A. Rahman, W. Olesinski, and P. Gburzynski. Controlled flooding in wireless ad-hoc networks. In *Proceedings of IWWAN'04*, Oulu, Finland, jun 2004.

[15] I. Stojmenovic, M. Seddigh, and J. Zunic. Dominating sets and neighbour elimination based broadcasting algorithms in wireless networksing protocol for wireless networks. *IEEE Transactions on Parallel and Distributed Systems*, 13(1):14–25, January 2002.

[16] M. T. Sun, W. C. Feng, , and T. H. Lai. Location aided broadcast in wireless ad hoc networks. In *Proceedings of GLOBECOM'01*, 2001.

[17] J. Wu and F. Dai. Broadcasting in ad hoc networks based on self-pruning. In *Proc. of INFOCOM*, March 2003.